

Every position has a level too.

Every US equity has a market benchmark — SPY does most of the work. Many have a useful sector benchmark — XLK explains a software stock that SPY alone doesn't. Some have a meaningful subsector benchmark — SOXX explains a chip stock that XLK alone doesn't.

But not every stock has all three.

Apple's subsector is so concentrated that the "subsector hedge" for AAPL is basically a noisy short of itself. Subtracting it makes the residual *worse*, not cleaner. The model knows this. It can tell you.

We call the answer **L***.

The picker

For each stock, on each trading day, RiskModels asks: how much does the next level down actually add?

```
if subsector adds ≥ 1% marginal variance explained:    L3
else if sector adds ≥ 1% marginal variance explained:   L2
else:                                                    L1
```

- **L1** — one ETF, the market. The default when sector and subsector don't earn their seat.
- **L2** — two ETFs, market + sector. The right answer for most large-cap stocks.
- **L3** — three ETFs, market + sector + subsector. Reserved for names where the subsector layer carries real signal.

The 1% bar isn't arbitrary. It came out of a 275-ticker × 11-sector study across two decades of returns. Tighter and the rule over-fits to noise. Looser and the rule drops to L1 too often.

What it changes for your workflow

| Old question | The L* answer |
|---------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| "What's the residual return on this name?" | <code>lstar_rr</code> — the residual at the level the model actually picked. Comparable across the whole universe. |
| "Should I hedge with three ETFs or two?" | <code>lstar_level</code> — 1, 2, or 3. Don't trade legs the model didn't prescribe. |
| "Why does my deep-hedge basket leak alpha on certain stocks?" | Because you were forcing L3 on names where L1 or L2 was the right depth. The subsector hedge added cost without adding cleanness. |

What clients see

`lstar_rr` and `lstar_level` ship as standard fields on every `/api/metrics` and `/api/batch/analyze` call — no separate endpoint, no extra cost. The cascade does the picking; you get the answer.

For a custom threshold (more aggressive or more conservative than the canonical 1%), the picker `/api/lstar?threshold=` endpoint gives you the full historical dispatch with your choice of bar.

```
from riskmodels import RiskModelsClient
client = RiskModelsClient.from_env()

m = client.get_metrics("AAPL")
m["lstar_rr"]      # → today's idiosyncratic return at AAPL's right depth
m["lstar_level"]  # → 3 (L3) for AAPL — subsector signal clears the bar
```

Where it matters

- **Stat-arb desks** ranking residual signals across a universe. `lstar_rr` is the cross-sectionally comparable residual — fixed-L3 isn't.
- **Allocators** sizing single-name risk. The `lstar_level` tells you whether the residual you're measuring assumed a one-, two-, or three-ETF hedge stack.
- **AI agents** answering "what's the residual?". L* keeps them from reaching for the deepest hedge by reflex when shallower is more honest.

The principle

| Hierarchical risk models give you more knobs. L^* tells you which knobs to actually turn.

Every position has a market benchmark. Every position has a level too.

[Try `lstar_rr` on the API →](#)